



ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



Sponsored by





ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



Sponsored by



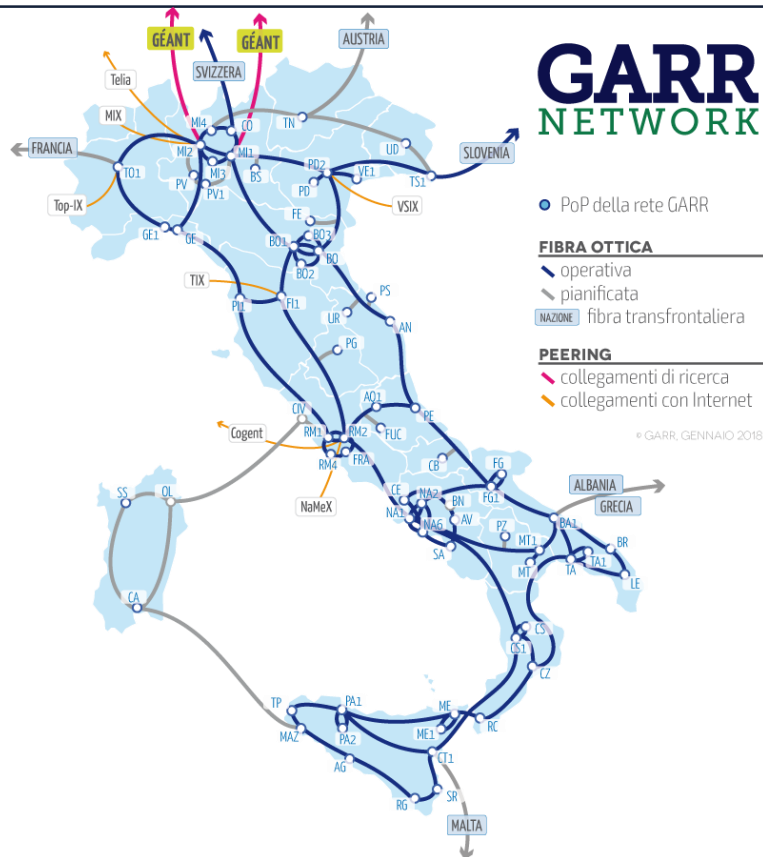
No Check Left Behind

Claudio Pisa - GARR

Rome
OpenInfraDays 2019



Consortium GARR



- Italian NREN (National Research and Education Network)
 - nonprofit organization
 - connected to GÉANT
- High speed network
 - 1000+ locations
 - ~4.5 million users
- teachers, researchers, students
 - 15,000 km of optical fiber
 - 380 Petabyte Yearly Traffic


GARR Computing and Storage

INFRASTRUTTURA DI CALCOLO E STORAGE DISTRIBUITO

📍 5 siti distribuiti

📊 8.448 virtual CPU

💻 10 PB spazio storage





ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



Sponsored by





- Tool for the deployment, configuration and management of services on public and private cloud infrastructures
- Free and open source
- Exposes a high-level declarative language
- **Charms**
 - script collections
 - can be written in any scripting/programming language
 - collected on the Juju store
- Juju orchestrates the deployment, composition and scaling of Charms

Monitoring - day 0 (one year ago)

- Different monitoring systems:
 - Nagios
 - OpenStack services
 - Zabbix
 - Ceph
 - Hardware sensors
- Some systems not monitored



ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



AGID | Agenzia per
l'Italia Digitale

Sponsored by



MESOSPHERE





ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



AGID | Agenzia per
l'Italia Digitale

Sponsored by

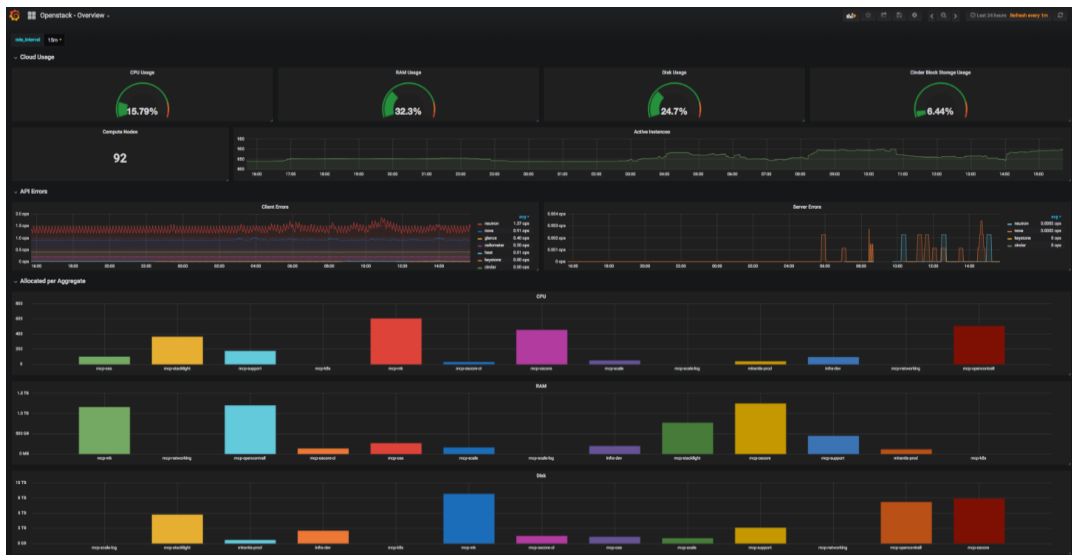


Objective

Objective:
**build a unified
comprehensive
dashboard**



State of the art:
Grafana



Considerations

- Existing monitoring systems seem to be doing their job well
 - the right tool for the right job
- What is missing is just a single viewpoint
- Grafana seems to have what we need



- **Grafana** is a platform for data visualization, querying and alerting
- Several pluggable data sources:
 - Zabbix
 - PNP (Nagios)
 - Prometheus
 - Gnocchi
 - Monasca
 - JSON (general purpose)
 - MySQL / PostgreSQL (general purpose)
- Data from heterogeneous sources can be mixed in the same dashboard



ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



Sponsored by




MESOSPHERE




- Zabbix
 - "batteries included" free and open source monitoring tool
 - auto discovery
 - XML based templates
 - remote agents
 - Good Ceph integration
 - Good multiuser support
- Straightforward integration with Grafana
 - Grafana's Zabbix datasource


Ceph PA




Ceph CT




Ceph PA sectors written




Ceph CT sectors written



Ceph PA sectors read



Ceph CT sectors read



Nagios

- **Nagios** is a free and open source monitoring tool
 - plugins - scripts to check and report
 - Nagios remote plugin executor (NRPE) - for remote hosts
 - known especially for alerting
 - FAQ: how do you pronounce Nagios?
 - the author pronounces it as “nah-ghee-ose”
 - but “you can pronounce it however the heck you'd like”
 - even “nachos”



Nagios

- Nagios is very well integrated in Juju
 - Nagios charm
 - NRPE charm
 - Canonical OpenStack charms come with handy Nagios configuration options

PNP my Nagios

- PNP is an addon to Nagios which analyzes performance data provided by Nagios plugins and stores them automatically into Round Robin Databases (RRD)
 - and there is a PNP Grafana datasource











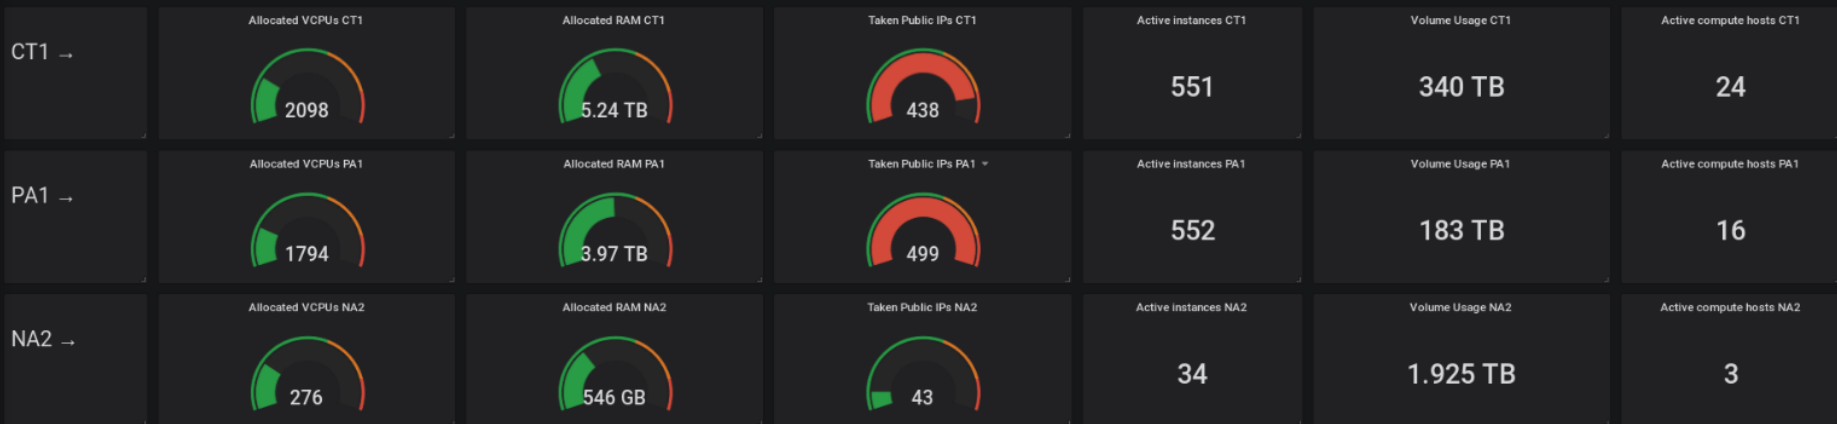
Using Nagios
for alerts



Using Nagios
for metrics



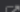






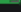
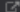



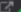
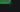


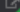
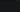





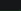
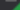
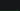

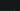


































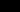
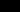
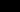
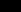
OpenStack → Nagios

- The Nagios Charm monitors the OpenStack services
 - but not what is happening inside OpenStack
- Simple idea: **write a set of Nagios plugins which collect metrics from the OpenStack API**
 - number of projects
 - number of servers
 - floating IP address usage
 - volume usage
 - OpenStack APIs reachability
 - ...



hosts

All ▾

 daas-pa-juju-controller-0 OK 	 daas-pa-juju-controller-1 OK 	 daas-pa-juju-controller-2 OK 	 juju-ceph-radosgw-ct1-cl1-19 OK 
 juju-ceph-radosgw-ct1-cl1-20 OK 	 juju-ceph-radosgw-ct1-cl1-21 OK 	 juju-cinder-ct1-cl1-47 OK 	 juju-cinder-ct1-cl1-48 OK 
 juju-cinder-ct1-cl1-49 OK 	 juju-cinder-pa1-cl1-4 OK 	 juju-cinder-pa1-cl1-5 OK 	 juju-cinder-pa1-cl1-7 OK 
 juju-controller-k8s-cloudmaster-ba OK 	 juju-controller-k8s-cloudmaster-ct OK 	 juju-controller-k8s-cloudmaster-pa OK 	 juju-controller-regions-ba OK 
 juju-controller-regions-ct2 OK 	 juju-ctrl-host-ct1-cl1-1 OK 	 juju-ctrl-host-ct1-cl1-2 OK 	 juju-ctrl-host-ct1-cl1-3 OK 
 juju-ctrl-host-pa1-cl1-0 OK 	 juju-ctrl-host-pa1-cl1-1 OK 	 juju-ctrl-host-pa1-cl1-2 OK 	 juju-glance-ct1-cl1-21 OK 
 juju-glance-ct1-cl1-22 OK 	 juju-glance-ct1-cl1-30 OK 	 juju-keystone-ct1-cl1-49 OK 	 juju-keystone-ct1-cl1-50 OK 
 juju-keystone-ct1-cl1-51 OK 	 juju-memcached-ct1-cl1-19 OK 	 juju-memcached-ct1-cl1-21 OK 	 juju-memcached-ct1-cl1-22 OK 
 juju-memcached-pa1-cl1-3 OK	 juju-memcached-pa1-cl1-4 OK	 juju-memcached-pa1-cl1-5 OK	 juju-nagios-server-pa1-cl1-0 OK

Kubernetes → Prometheus → Grafana

- A new kid on the block: Kubernetes
 - container platform
 - inspired by **Google Borg**
- Prometheus
 - open source monitoring tool
 - inspired by the **Google Borg Monitor**
 - powerful query language (PromQL)
 - alerting
 - white-box monitoring
- Kubernetes supports Prometheus natively
- Grafana supports Prometheus natively





ITALY
OpenInfra Days



Organized by

IRIDEOS



Under the patronage of



Sponsored by



Instance All Namespace kube-system

all pods

Memory Working Set

Last 1 minute



Cpu Usage

Last 1 minute



Filesystem Usage

Last 1 minute



Used

49.17 GiB

Total

1.60 TiB

Used

3.29 cores

Total

224.00 cores

Used

529.09 GiB

Total


2.39 TiB


Network



each pod

Cpu Usage





Considerations / Takeaways

- Nagios (vs. Prometheus)
 - Nagios has no query language
 - but it is very easy to develop new plugins
 - Nagios (+PNP) uses RRD based storage
 - not suitable for highly dynamic environments
 - e.g. cloud-native applications
 - but suitable for infrastructure monitoring
- Grafana performance (time to render graphs)
 - very good with Nagios+PNP
 - OK with Zabbix
 - can be slow with Prometheus

Future Work

- Self healing
 - react to well known bugs
 - with well known recipes
 - with some hysteresis / guard time
- AIOps - Artificial Intelligence for Operations
 - collect many many metrics
 - annotate incidents/events
 - train an AI using the collected data
 - profit!



Thank you

Claudio Pisa - GARR - Distributed Computing and Storage Department

claudio.pisa@garr.it